ingress and egress node. It is determined if a predetermined time out period has expired since detection of the failure. If the time out period has expired, it is determined if the failure has been corrected. If the failure has not been corrected, resources associated with the permanent sub network connection are deallocated.

Advantageous implementations can include one or more of the following features. Determining if a predetermined time out period has expired can include retrieving a time out period value associated with the failed permanent sub-network connection and initiating a timer with the time out period value. Deallocating resources can include signaling, by one or more nodes in a path forming the permanent sub-network connection between the ingress and egress nodes, to other nodes in the path instructions to tear down the path. Route information associated with the permanent sub-network connection can be stored prior to tear down such that at a time for restoring the permanent sub-network connection, no optimal routing determination is required.

In another aspect, the invention provides methods and apparatus, including computer program products, implementing and using techniques for deallocating resources in a network of connected nodes. A failure in a path in the network is detected. It is determined if the path includes a permanent sub-network connection, and if so, for each permanent sub network connection it is determined if a predetermined time out period has expired since detection of the failure. If the time out period has expired, it is determined if the failure has been corrected. If the failure has not been corrected, resources associated with the permanent sub network connection are deallocated.

In another aspect, the invention provides methods and apparatus, including computer program products, implementing and using techniques for deallocating resources in a network of connected nodes. A failure in a path in the network is detected. Resources for all sub-network connections traversing the path are immediately cleared. A pre-determined time out period is waited prior to clearing all resources for each permanent sub-network connection traversing the path.

The invention can be implemented to realize one or more of the following advantages. Connection routes can be set up automatically between ports coupled to the nodes of the network. There may be no need for a system administrator to manually program the route in each node in the network between the ingress point and the egress point. Upon

provided through a signaling and routing protocol (e.g., OSRP). Explicitly provisioned SNCs include user (e.g., system administrator) -defined routes. Automatically provisioned SNCs make use of a routing protocol (e.g., as implemented in routing unit 250) for computing an optimal route. In either case, the route information is transmitted to other nodes in the network and cross-connects associated with the routes are configured. The SNCs are said to be temporary in that, resources associated with the route (e.g., bandwidth) can be dynamically re-allocated along the path. The reconfiguration includes the clearing of the set up connection (e.g., freeing the resources at a given node). Network resources associated with the SNCs are dynamically reconfigurable. Accordingly, the failure at a single point along the path from an ingress node to an egress node defining the route will not result in unused and unavailable resources. In one implementation, a user can configure one or more of the following parameters associated with a SNC including a local line on which the SNC originates, the identification (ID) of the node on which the SNC terminates, the ID of the remote line on which the SNC terminates, a class of service, a maximum allowable delay, route setting including working and protection routes, preferred status, mesh restorability, revert configurations upon failover and reversion timers.

A permanent SNC (P-SNC) defines a hybrid of a permanent and sub-network connection. A P-SNC is characterized as being permanent in that, upon failure, the SNC is not torn down immediately and the resources associated therewith are not deallocated. However, the P-SNC can be set up and torn down with signaling. A timer (e.g., timer block 251 in signaling unit 250) is associated with each P-SNC. The timer defines an amount of time that may expire prior to the tear down of a given connection after a failure. Accordingly, a P-SNC is temporary in that, after the timer expires, the connection can be torn down and resources reallocated as appropriate. A P-SNC entry in the path specification table 260 can include a field that specifies the time out associated with tear down of the given P-SNC. Signaling unit 250 can initialize, at the detection of a failure along the P-SNC connection, a timer with a value as specified in the path specification table for a given P-SNC connection. After the time out has expired, signaling unit 250 can tear down the route and de-allocate the resources dynamically without requiring the manual reconfiguration in nodes along the path defined by the route.